

When a missile strike can knock out cloud infrastructure across a region, protecting data centres must become as central to national security as air defence, energy grids and nuclear command-and-control.

At approximately 4:30 on a Sunday afternoon in Dubai, **emergency crews cut power to an Amazon Web Services data centre** after objects struck the facility, igniting a fire. Connectivity across that availability zone went dark. Engineers scrambled to reroute traffic through neighbouring regions. The outage, AWS warned, would take several hours to resolve.

The timing was not coincidental. **Iranian missiles and drones were, at that same moment, raining down on targets across the Gulf**, Qatar, Kuwait, Saudi Arabia, the UAE, in retaliation for the coordinated U.S.-Israeli strikes, codenamed Operation Epic Fury and Roaring Lion, that had begun four days earlier on February 28th, killing Supreme Leader Ayatollah Khamenei and striking nuclear sites at Isfahan, Qom, Karaj and elsewhere across Iran. The physical and digital battlespaces had, in one afternoon, fused into a single theatre of war.

AWS has not formally confirmed the link between the projectile strikes and its UAE data centre fire. It does not need to. The image alone, a smoking server hall, dark screens, rerouted packets, tells us something that Western defence establishments have been slow to absorb: the data centre is now a frontline military asset, as strategically vital as the aircraft carrier, the missile silo, or the power grid. And we are nowhere near ready to protect it like one.

If a seven-hour peacetime software failure can quietly ground thousands of planes and suspend surgical schedules across a continent, what happens when an adversary does it on purpose?

The Stack We Forgot to Defend

The dependency is easier to see once you look for it. Think of it as a three-layer stack. At the bottom sits physical infrastructure, fibre runs, undersea cables, power substations, and the anonymous buildings ringed by chain-link fences that house the world's compute. Above that layer sit the hyperscale clouds: Amazon Web Services, Microsoft Azure, Google Cloud. And riding on top of those clouds, entirely invisible to most citizens, are the systems that run daily life, airline reservation platforms, hospital electronic health records, payment processing engines, government identity and benefit systems, and the targeting algorithms, logistics networks and secure communications of modern militaries.

The fragility of that stack becomes undeniable the moment any layer fails. In July 2024, a **botched software update pushed to a widely deployed security agent caused more than five thousand flight cancellations worldwide in a single day**. Airlines' check-in systems, crew-scheduling platforms and reservation engines simply stopped working. **A single point of failure, in a single software component, paralysed global aviation for hours**. That same year, a **seven-hour outage at AWS disrupted electronic health records**, telemedicine platforms and hospital billing systems across the country, forcing clinicians back to paper charts, delaying surgeries, and costing healthcare providers an estimated sixty-two thousand dollars per hour in lost productivity. No bomb was dropped. No adversary acted. The system simply hiccupped, and the consequences cascaded.

Now ask the harder question: if passive failure produces this much damage, what does deliberate attack produce? The answer is not a thought experiment. It is already being written in Ukraine.

AcidRain and the Opening Salvo

In the early hours of February 24th, 2022, before Russian tanks crossed the Ukrainian border, Russian military intelligence launched a cyberattack against the ground infrastructure supporting Viasat's KA-SAT satellite internet service. A piece of **malware later named AcidRain wiped the firmware of tens of thousands of satellite modems in Ukraine and across parts of Europe**. Ukrainian military units that depended on that service for satellite-based communications lost them in the opening minutes of the war, **complicating command-and-control at precisely the moment it mattered most**. The collateral damage spread even further: wind turbines in Germany, operated remotely through modems on the same network, went offline. An energy-sector disruption, in Germany, from a cyberattack timed to the first hours of a land war in Ukraine.

The lesson that military planners drew , or should have drawn , is this: digital infrastructure attacks are no longer a side effect of modern warfare. They are the opening act. They are planned, rehearsed, and executed before the first kinetic strike, precisely because blinding an adversary's communications and command networks is worth more, in the critical opening hours, than destroying any single weapons platform.

The current conflict in the Gulf is confirming the lesson in real time. Cyber-threat bulletins from late February 2026 document Iranian actors conducting denial-of-service attacks, probing industrial control systems, and staging network intrusions against regional infrastructure in parallel with the missile barrages. Over one hundred and fifty separate hacktivist incidents have already been recorded, with security analysts warning of spillover risks to energy, finance and IT sectors well beyond the immediate conflict zone. These are not nuisance attacks. They are invisible salvos, exchanged alongside the visible ones, with potentially strategic consequences that receive a fraction of the public scrutiny.

The Viasat attack did not make headlines the way a tank column does. But it shaped the opening hours of the largest land war in Europe since 1945.

The Military's Hidden Dependency

The vulnerability is not confined to the civilian stack. Modern militaries have become profoundly dependent on data infrastructure in ways that are only beginning to be acknowledged publicly. Intelligence, surveillance and reconnaissance feeds from satellites, drones and sensors are ingested and fused in data centres. Targeting algorithms run on servers. Logistics, maintenance and personnel systems, the unglamorous connective tissue of a fighting force, live in the cloud. Secure communications depend on the availability of computing infrastructure that sits somewhere, physically, in a building that can be struck.

The drive toward edge computing and battlefield AI has, if anything, sharpened the dependency. Military planners now want micro-data centres deployed close to the front, reducing the latency of autonomous systems and drone operations to microseconds. In high-tempo warfare, those microseconds matter. Which means the compute node itself, small, distributed, often commercially sourced, has become a tactical asset whose protection rivals that of any weapons platform. Losing a targeting server at the wrong moment does not just degrade capability. It can alter the outcome of an engagement.

The defence industry has quietly absorbed this. Contractors building secure data centres for military clients design to specifications that include resistance against bombardment, electromagnetic pulse shielding to prevent signal interception, and energy resilience engineered to survive grid disruption. The facilities are, in every meaningful sense, hardened

military installations , just ones that happen to look, from the outside, like anonymous warehouses on a business park.

The problem is that civilian infrastructure , the hyperscale cloud platforms on which vast swathes of both military logistics and civilian essential services now depend , is built to none of those standards. It is built for uptime and cost efficiency, not for survival in a contested electromagnetic and kinetic environment.

The Policy Gap We Cannot Afford

Regulators and alliance structures have begun to grasp the problem, but their responses remain fragmented and largely reactive. NATO formally recognised cyberspace as an operational domain in 2016 and now operates a Cyber Security Centre and a Cyberspace Operations Centre providing around-the-clock defence of alliance networks. The European Union's NIS2 Directive and its Critical Entities Resilience Directive impose cybersecurity and continuity obligations on operators of essential services, including energy, transport, banking and digital infrastructure. These are real steps. They are not enough.

The core vulnerability that neither framework adequately addresses is the concentration of mission-critical dependencies in a handful of hyperscale commercial cloud providers. The AWS healthcare outage of 2024 exposed the depth of the problem: hospitals across the country had, often without fully realising it, centralised their most sensitive and time-critical systems , patient monitoring, electronic records, surgical scheduling , with a single vendor. When that vendor's networking infrastructure failed, everything failed simultaneously. There was no fallback, because no one had been required to build one.

In wartime, a sophisticated adversary will not wait for a vendor to have a bad day. It will find those single points of failure and exploit them deliberately , through cyberattack on control

planes, kinetic strikes on key data centre regions, **sabotage of the undersea cables that carry 99% of international internet traffic**, or electromagnetic attack designed to degrade compute at scale. The choke points are not hidden. They are, for anyone who cares to look, astonishingly legible.

Analysts studying data centres as twenty-first-century battlespace have argued that computational capacity has become a determinant of geopolitical power in the same way that naval tonnage or oil reserves once were. Who controls hyperscale GPU-rich infrastructure leads in autonomous systems, AI-enabled intelligence, and offensive cyber operations. That framing is correct, and it implies a doctrine that does not yet exist. We need explicit deterrence thresholds for attacks on cloud infrastructure: statements of intent, backed by capability, that communicate to adversaries the consequences of targeting civilian digital assets that underpin military function. We need mandatory multi-cloud redundancy for essential services, automated failover between commercial and government clouds, and publicly funded resilience reserves. And we need the conversation about what constitutes critical national infrastructure to be updated, urgently, to reflect the world that actually exists in 2026.

We do not allow essential power grid components to be owned and operated with no resilience standards. We should not allow the digital infrastructure they now depend on to be treated any differently.

The Reckoning

The twentieth century's paradigm of critical infrastructure was built around things you could

see and touch: refineries, power stations, missile silos, carrier battle groups. The logic of deterrence and defence was organised around protecting those assets, because their loss meant strategic vulnerability. The twenty-first century has added a new category , one that is equally determinative, equally vulnerable, and almost entirely absent from the public conversation about national security.

The data centre humming behind an anonymous fence in the Dubai desert was not, until last Sunday afternoon, anyone's idea of a frontline installation. It is now. The fire that broke out when projectiles struck it did not just disrupt cloud services for a few hours. It demonstrated, in the most visceral possible terms, that the boundary between the digital and physical dimensions of warfare has collapsed , and that we have built an extraordinary amount of our military capability, our economic function, and our civilian essential services on infrastructure that was never designed to survive a war.

In the Gulf right now, missiles and packets are flying simultaneously toward the same strategic objectives. When this conflict ends , and all conflicts end , the question that its architects and adversaries will study is not just which weapons systems performed. It will be which data centres stayed up.

We should start asking that question before the next one begins.